Evolving Neural Network Agents to Play Atari Games with Compact State Representations

Adam Tupper **Kourosh Neshatian**

Department of Computer Science and Software Engineering, University of Canterbury, New Zealand

Overview

- **The size of deep neural networks that have shown success in solving complex reinforcement learning (RL) problems limits the effectiveness and benefits of neuroevolution** methods that have proven effective at solving simpler RL problems in the past.
- **A potential solution to this problem is to separate state representation and policy learning**, and only apply neuroevolution to the latter.
- We extend research following this approach [1, 4] by **evolving small policy networks for Atari games using NEAT** [5], that learn from compact state representations provided by the recently released Atari Annotated RAM Interface (AtariARI) [2].

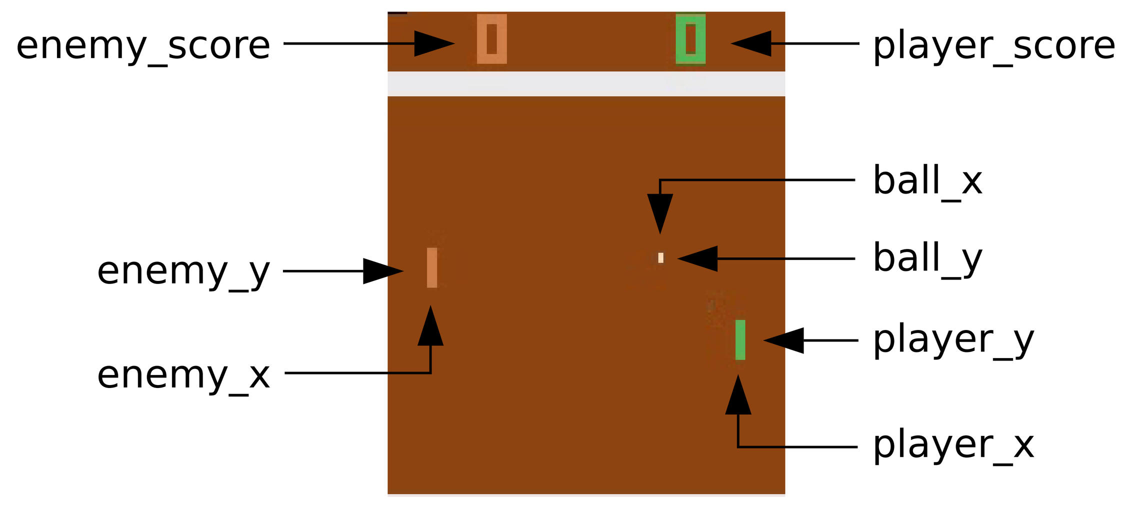

Method

- **We used NEAT** to evolve agents that take, as input, the values of state variables provided by the AtariARI.
- **The AtariARI identifies the bytes of RAM that store important state variables**, reducing the size of the input space by up to 93% when compared to using the entire contents of RAM.
- **We chose 14 games to use in our evaluations**, based on our assessment of the perceived completeness of the information provided by the AtariARI. These games have representations that are either *good* (complete) or *fair* (near-complete).
- **A single set of hyperparameter values were used to evolve a separate agent for each game**. These were chosen based on informal experimentation on a subset of three games: Asteroids, Boxing and Pong.
- **For each game, three evolutionary runs were performed**, each lasting 200 generations, with a population size of 130. The reported results are for the best performing agents among those runs.
- **The performance of each agent is compared to expert human scores** published alongside DQN [3].

Overall Performance

The agents for Video Pinball, Boxing, Bowling, and Freeway exceed or are competitive against expert human performance.

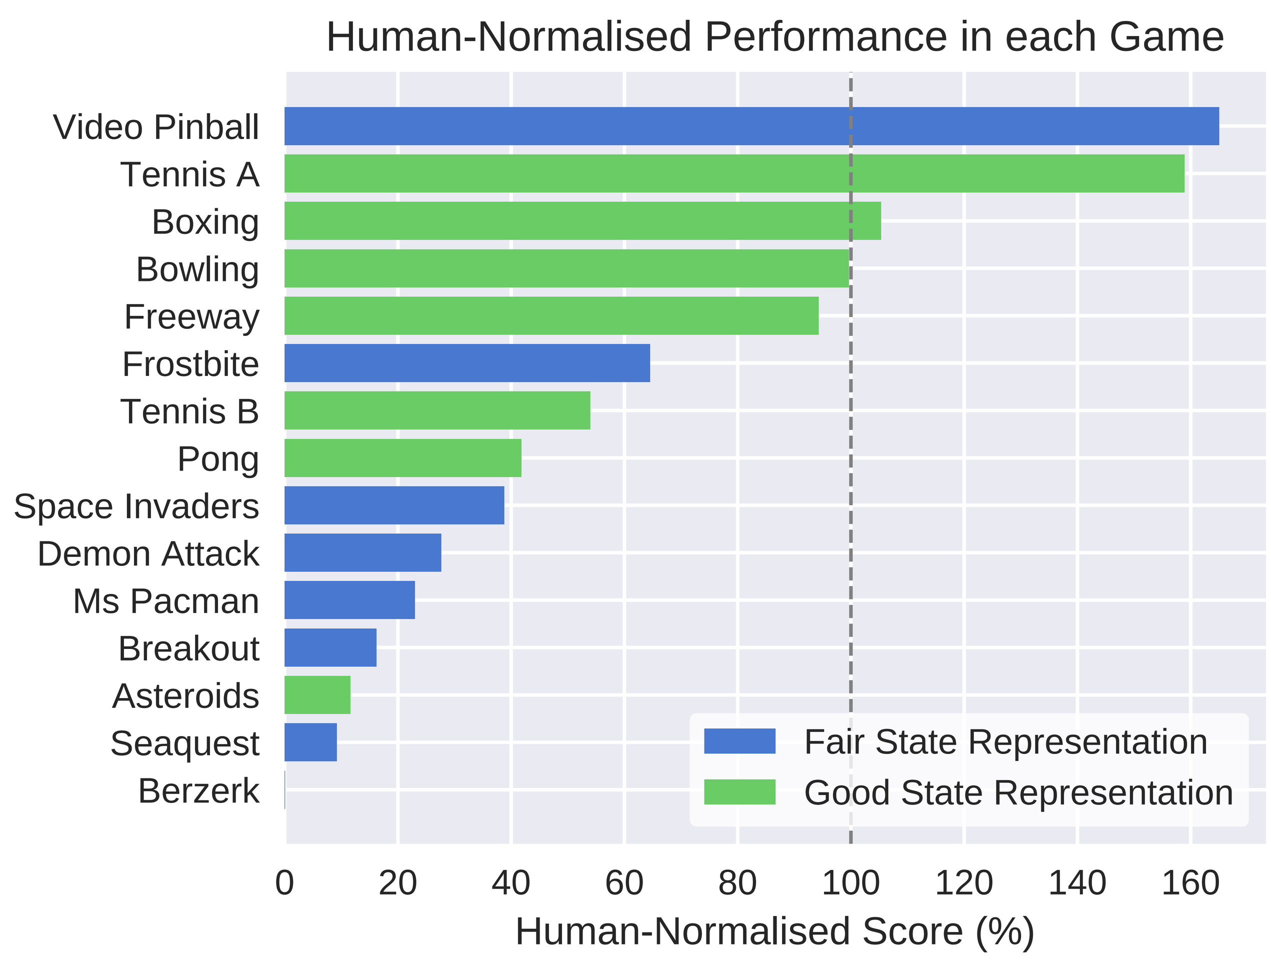

Although most of the best performing games have *good* state representations, the highest performing agent was found for Video Pinball. This illustrates that for some games, **good strategies can still be discovered with imperfect information**.

Two conditions are reported for Tennis because initially (Tennis A), the evolved agent exploited a loop-hole in our setup by refusing to serve the ball, instead waiting for the episode frame cap to limit it's losses.

Evolved Architectures

A surprising aspect of our results was the **simplicity of the evolved architectures for some high-performing solutions**. The best-performing agents for Freeway and Bowling epitomise this simplicity, utilising **only one, and no hidden nodes** respectively.

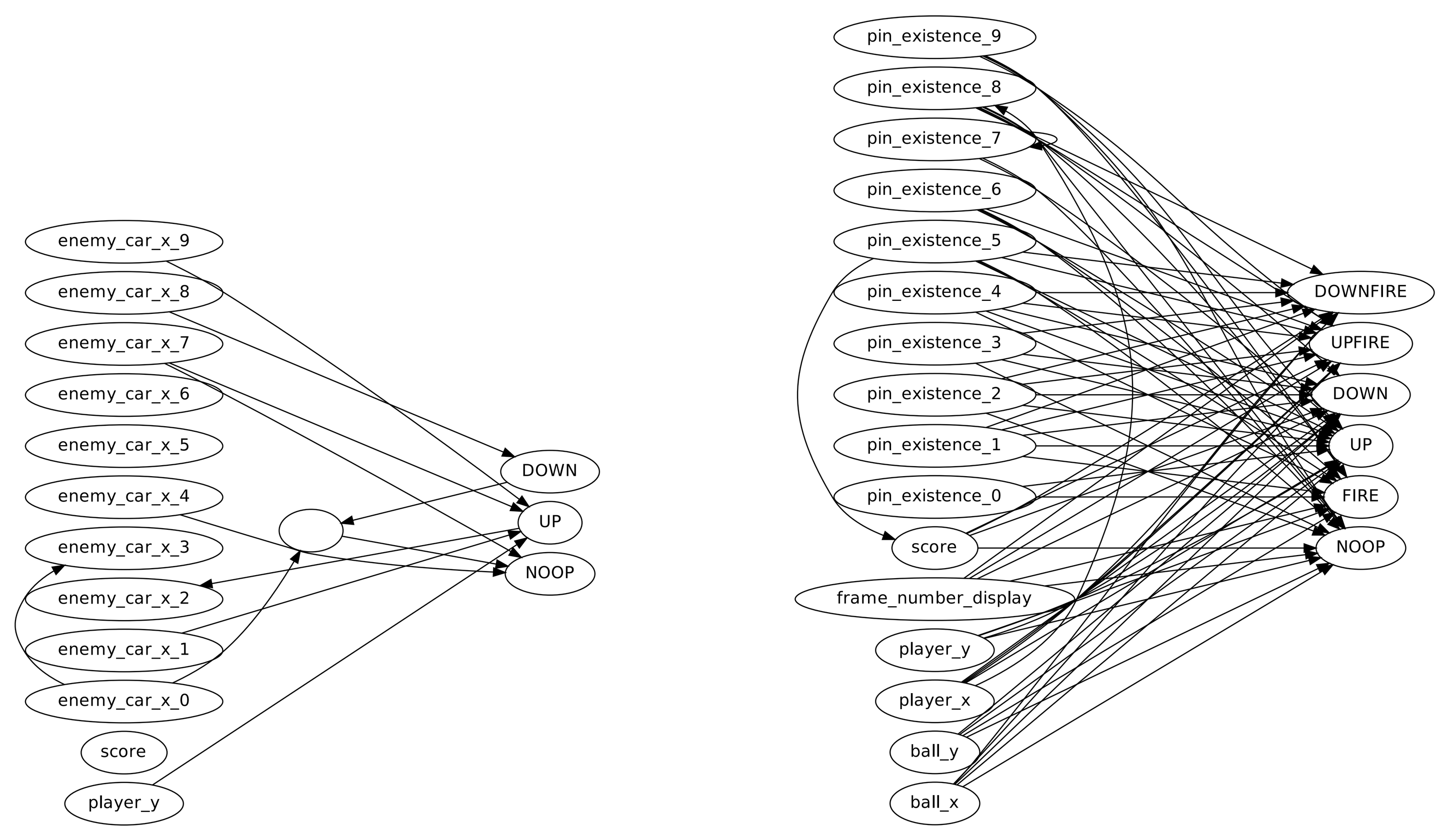

(a) Freeway agent (b) Bowling agent

Although the simplicity of some Atari games is widely known, these results show just how simple some solutions can be. Training a fixed architecture would likely never lead to such solutions, highlighting the benefit of TWEANN neuroevolution methods.

Conclusions & Future Work

- Although evolved policies only exceeded or were competitive with expert human performance in a handful of games, we discovered that surprisingly simple and small neural networks could play these game effectively.
- In our ongoing work, we are investigating the performance of a separated state representation and policy learning framework that uses NEAT to evolve policy learners from *learned* compact state representations.

References

[1] Alvernaz, S. and Togelius, J. 2017. Autoencoder-augmented neuroevolution for visual doom playing. 2017 IEEE Conference on Computational Intelligence and Games (CIG), 1–8.
[2] Anand, A. et al. 2019. Unsupervised State Representation Learning in Atari. Advances in Neural Information Processing Systems 32, 8766–8779.
[3] Mnih, V. et al. 2015. Human-level control through deep reinforcement learning. Nature. 518, 529.
[4] Poulsen, A.P. et al. 2017. DLNE: A hybridization of deep learning and neuroevolution for visual control. 2017 IEEE Conference on Computational Intelligence and Games (CIG), 256–263.
[5] Stanley, K.O. and Miikkulainen, R. 2002. Evolving Neural Networks through Augmenting Topologies. Evolutionary Computation. 10, 99–127.

More Information

evolvingatari.adamtupper.nz